JMLR: Workshop and Conference Proceedings 60:1–10, 2016 ACML 2016

Reinforcement Learning Approach for Parallelization in Filters Aggregation Based Feature Selection Algorithms

Ivan Smetannikov ISMETANNIKOV@CORP.IFMO.RU
Ilya Isaev ISAEV@RAIN.IFMO.RU
Andrey Filchenkov AFILCHENKOV@CORP.IFMO.RU
ITMO University, St. Petersburg, Kronverksky Pr. 49

arXiv:1611.02047v1 [cs.LG] 7 Nov 2016

Abstract

One of the classical problems in machine learning and data mining is feature selection. A feature selection algorithm is expected to be quick, and at the same time it should show high performance. MeLiF algorithm effectively solves this problem using ensembles of ranking filters. This article describes two different ways to improve MeLiF algorithm performance with parallelization. Experiments show that proposed schemes significantly improves algorithm performance and increase feature selection quality.

Keywords: Machine learning, feature selection, rank aggregation, multi-armed bandit, parallel computation, MeLiF, MeLiF+, PQMeLiF, MAMeLiF.

1. Introduction

Almost all business and scientific problems nowadays involve processing huge amounts of data with machine learning algorithms. Due to its universal applicability, machine learning became one of the most promising and researched scientific domains. In particular, it has application in bioinformatics (Bolón-Canedo et al., 2014; Saeys et al., 2007), as giant amounts of data about gene expression of different organisms are obtained in this field. In order to filter data noise and reduce model complexity, it is necessary to select the most relevant features. Techniques and methods achieving this goal are called feature selection. Gene expression data can enable researchers to spot which DNA pieces are responsible for reactions to particular environment change or some internal processes of an organism. The main problem met in processing such data is the high dimensionality of instances. Gene expression datasets often have a high number of features and relatively low number of objects. For a dataset with these properties, it is very hard to build a model that fits the data well.

A feature selection algorithm meets several requirements. It is expected to work fast and show good performance. However, no universal algorithm for feature selection exists. Wrappers (Kohavi and John, 1997) are the family of methods based on searching for an optimal feature subset that maximizes preselected classifier effectiveness. Such problem statement leads to high performance of a found solution. However, the size of search space grows exponentially of the instance dimensionality. This fact makes wrapper rarely applicable in bioinformatics, as the number of features in datasets could be up to hundreds of thousands. In these cases, other feature selection algorithms known as filters (Sánchez-Maroño et al., 2007) are used. Filters are based on estimation of feature importance. Filters usually

© 2016 I. Smetannikov, I. Isaev & A. Filchenkov.

perform worse than wrappers, but they are much faster. A special group of feature selection methods are embedded selectors (Lal et al., 2006) that uses particular properties of a selected classifier.

Ensembling, which is the process of building a combination of several simple algorithms, is a widely used technique in machine learning (Bolón-Canedo et al., 2012). MeLiF algorithm proposed in (Smetannikov and Filchenkov, 2016), applies ensembling to feature selection. This algorithm tries to find such linear combination of basic ranking filters, which selects the most relevant features of the dataset. Ranking filter itself consists of two separate parts: a feature importance measure and a cutting rule. Basically, MeLiF tries to tune coefficients of feature importance measure linear combination. This process involves classifier training, evaluating and comparing with ranking filters themselves, thus making it comparatively slow. This is why parallelization can become really handy and helpful to improve algorithm computational time.

The simplest parallelization scheme called MeLiF+ is described in (Isaev and Smetannikov, 2016). The main disadvantage of this naïve scheme is that it does not scale well. It starts search of the best coefficient vector from several starting points using a separate thread for each point. When one of these optimization processes ends, this thread just stops and its resources stay unreleased therefore they cannot be used for further work. Thus, it is not useful to allocate a lot of resources for this process as most of them will stay unused.

To overcome this problem, it is necessary to use cores of processing server more effectively. While processing, MeLiF visits a lot of points in the linear space, so we can process points using a task executor. This research proposes two different approaches to using parallel coordinate descent in building ensembles of ranking filters called PQMeLiF and MAMeLiF. The first algorithm stores points that should be processed in a priority queue. The second algorithm solves theparallelization problem by reducing it to the multi-armed bandit problem.

The remainder of the paper is organized as follows: Section 2 describes MeLiF algorithm, Section 3 contains the proposed parallelization schemes, Section 4 outlines experimental setup, Section 5 contains experiment results, and finally Section 6 contains conclusion.

This paper is a version of the paper accepted to 5th International Young Scientists Conference in HPC and Simulation.

2. Linear combination of ranking filters

Ranking filter f is a pair $\langle m, \kappa \rangle$, where m is a feature importance measure and κ is a cutting rule. For each object feature, m return its importance for label prediction. For a sorted list of features, κ cuts the worst. The core idea of MeLiF is to use several ranking filters $f_1, \ldots, f_N$ in order to merge them into a single ranking filter by finding the most effective linear combination of their feature importance measures. This combination is a new feature importance measure, while the cutting rule can be inherited from all the ranking filters (usually, it is chosen empirically). Any performance measure may be used to evaluate a ranking filter effectiveness. In this paper, we use classifier effectiveness estimated with F_1 score.

Thus, MeLiF simply optimizes a function in the N-dimensional space, where N is the number of basic ranking filers. Evaluation of this function is comparatively costly (we need

to run a classifier). However, dimensionality of the search space is less by an order of magnitude than the number of features. This detail allows classifying an algorithm from MeLiF family as a hybrid of filter and wrapper, which inherits filter speed and wrapper focus on resulting performance.

The algorithm is parametrized with the following hyperparameters:

- $\delta \in \mathbb{R}$, a value of grid spacing;
- $P \in 2^N$, starting points;
- *evaluate*, the function for classifier effectiveness evaluation at given point in the search space.

The original MeLiF performs coordinate descent in the search space. It has been observed during experiments that the best option is this particular choice of starting points: $(1, 0, \ldots, 0)$, $(0, 1, \ldots, 0), \ldots, (0, 0, \ldots, 1)$ corresponding to the only one basic ranking filter used, and $(1, 1, \ldots, 1)$ corresponding to the equally weighted combination of all the basic ranking filters.

For each point it reaches, MeLiF tries to shift each coordinate value to $+\delta$ and δ. Then it evaluates effectiveness of each point. If evaluation result is greater than the current maximum, the algorithm assigns the current maximum to be equal to the coordinates of this point and starts searching from its first coordinate. If all coordinates are shifted to $+\delta$ and $-\delta$ and no quality improvement is observed, then the algorithm stops.

For each point obtained during the coordinate descent, the algorithm measures the value of the resulting linear combination of basic filters for each feature in the dataset. After that, the results are sorted, and the algorithm selects m topmost features. They are used to train and test a particular classifier. The classification quality is treated as point score. It is cached and compared to other points.

3. MeLiF parallel optimizations

In this paper, we propose two parallelization schemes of MeLiF and show that some of their configurations have speed improvement growing linearly of the number of processors. Furthermore, the proposed schemes show equal or even better performance quality in comparison with the single-threaded version of MeLiF.

The first proposed parallelization scheme is named PQMeLiF. In this name, PQ stands for priority queue. This algorithm is a variation of the best-first search algorithm (Russell and Norvig, 2009). The algorithm stores points that should be processed in a priority queue. On each iteration, the algorithm polls a point from the queue, calculates its score, and puts all its not visited neighbors back to the queue with the priority equal to the calculated score. Before initiating the algorithm, starting points should be put into a queue with maximum priority, which is 1.0. This ensures that all of them will be processed at the very beginning, so all the starting points will be taken into account simultaneously.

Unlike MeLiF and MeLiF+, PQMeLiF enables to tune halting criteria to find a trade-off between feature selection quality and algorithm performance: we can limit the number of points that should be processed. Experiments determined an optimal number that enables the algorithm to perform better than the original MeLiF, but as fast as possible.

Algorithm 1 MeLiF pseudo code

Require: starting points, δ, *evaluate*
  $q^* = 0$
  p^*
  for $p : points$ **do**
    $q = evaluate(p)$
    if $q > q^*$ **then**
      $p^* = p$
      $q^* = q$
    end if
  end for
  $smthChanged = true$
  while $smthChanged$ **do**
    for $dim : p.size$ **do**
      $p+ = p\{p[dim] + \delta\}$
      $q+ = evaluate(p+)$
      if $q+ > q^*$ **then**
        $q^* = q+$
        $p^* = p+$
        $smthChanged = true$
        $break$
      end if
      $p- = p\{p[dim] - \delta\}$
      $q- = evaluate(p-)$
      if $q- > q^*$ **then**
        $q^* = q+$
        $p^* = p+$
        $smthChanged = true$
        $break$
      end if
    end for
  end while
  return p^*, q^*

Algorithm 2 PQMeLiF pseudo code

```
Require: starting points, δ, evaluate, T
  q = PriorityBlockingQueue
  p*
  for p : points do
    enqueue(q, p, 1.0)
  end for
  for each new thread run in T threads do
    while !mustStop() do
      p = dequeue(q)
      score = evaluate(p)
      updateBestScore(score, p)
      neighbours = getNeighbours(p, δ)
      for p2 : neighbours do
        enqueue(q, p2, score)
      end for
    end while
  end for
  return point*
```

The experiments described in the next two sections showed that PQMeLiF performed much better than MeLiF and MeLiF+. However, they uncovered its drawbacks. The algorithm starts working initially with only the set of starting points in the queue, number of which is fixed and equals to the number of the basic ranking filters plus one. So, if a server has more cores, the extra cores stay unused until new points are added to the queue. The possible solution for this problem is to process more starting points in order to keep all the server cores busy. This algorithms drawback can be overcome with our next algorithm.

The main idea of this algorithm is to consider the problem of selecting new points as a reinforcement learning problem, in which we need to find a trade-off between exploration (of new areas in the search space) and exploitation (by evaluating points in areas where we have already found many good points) (Sutton and Barto, 1998). In order to apply this idea, we adopted the well-known UCB1 algorithm (Auer et al., 2002) for our parallelization problem by reducing it to multi-armed bandit problem. This reduction was performed in the following way. Firstly, we split the search space into different areas and correspond each area with an arm. Then, evaluating a point in an area is understood as playing the corresponding arm. Reward obtained by such playing is the score of the evaluated point (Bubeck and Slivkins, 2012; Desautels et al., 2014).

The problem we faced during implementation of reinforcement learning algorithm in multi-agent environment is delayed feedback. Basically, when some thread needs to select which arm to use, there is no information about the results of other threads at that moment. So it is forced to make a decision with the lack of information. Authors of (Joulani et al., 2013) proposed to make a decision based only on results that were already computed and provide a theoretical proof, that an error function is additive and based on the duration of delay.

Algorithm 3 MAMeLiF pseudo code

Require: starting points, δ, *evaluate*, T
  $points = splitSearchSpace(P)$
  $queues = []$
  q^*
  p^*
  for $p : points$ **do**
    $q = PriorityBlockingQueue$
    $enqueue(q, p, 1.0)$
    $queues+ = q$
  end for
  for *each new thread run in* T *threads* **do**
    while !$mustStop()$ **do**
      $qNE = findNonEmpty(queues)$
      $q = findBestQ(qNE)$
      $p = dequeue(q)$
      $score = evaluate(p)$
      $updateBestScore(score, p)$
      $neighbours = getNeighbours(p, \delta)$
      for $p2 : neighbours$ **do**
        $enqueue(q, p2, score)$
      end for
    end while
  end for
  return $point^*$

4. Experimental setup

As a classifier, we used SVM with polynomial kernel and soft margin parameter $C = 1$ implemented in WEKA library. We used 5-fold cross-validation. The number of selected features was constant: $m = 100$.

We ran our experiments on a machine with the following characteristics: 32-core CPU AMD Opteron 6272 @ 2.1 GHz, 128 GB RAM. We used $K = 50$ threads, where $K = 2pf$, p is the number of starting points, and f is the number of folds.

As the basic filters, we used Spearman Rank Correlation, Symmetric Uncertainty, Fit Criterion, VDM (Auffarth et al., 2010). We also executed MeLiF and MeLiF+ and recorded work time and point with the best classification result.

We used 36 datasets of different sizes from these archives: GEO, Broad institute. Cancer Program Data Sets, Kent Ridge Bio-Medical Dataset, Feature Selection Datasets at Arizona State University, RSCTC2010 Discovery Challenge All these datasets are DNA-microarray datasets with high number of features (from a few thousands up to a few dozens of thousands) and comparatively low number of objects (less than few hundreds).

5. Results

We used several different configurations of halt criteria in our algorithms that described below. For each dataset, we conducted 10 experiments total:

- B: basic one-thread MeLiF;
- P: naive parallelization method MeLiF+;
- PQ75: PQMeLiF limited to visit 75 points;
- PQ100: PQMeLiF limited to visit 100 points;
- PQ125: PQMeLiF limited to visit 125 points;
- PQrel: PQMeLiF that stops when no quality increase was registered in the previous 32 points;
- MA75: MAMeLiF limited to visit 75 points;
- MA100: MAMeLiF limited to visit 100 points;
- MA125: MAMeLiF limited to visit 125 points;
- MArel: MaMeLiF that stops when no quality increase was registered in the previous 32 points;

Each one of them halted if it gets the highest score 1.0. The results of the basic MeLiF, MeLiF+ are presented in the Table 1. Only best configurations for PQMeLiF and MAMeLiF are presented in it, that are PQ75 and MArel correspondingly. As it could be seen from the Table 1, PQMeLiF and MAMeLiF strongly outperform MeLiF+ resulting in approximately linear scaling over the number of computational cores. Also new methods gain in average small boost to the original MeLiF feature selection quality.

Table 1: Algorithms comparison

Dataset	Time				F_1 score			
	MeLiF	MeLiF+	PQ75	MArel	MeLiF	MeLiF+	PQ75	MArel
Arizona1	558	85	85	117	0.833	0.833	0.833	0.833
Arizona5	219	67	37	48	0.768	0.79	0.786	0.773
Breast	161	24	27	17	0.844	0.822	0.812	0.802
CNS	33	6	7	6	0.742	0.791	0.899	0.83
Data_train0	172	71	28	32	0.853	0.853	0.849	0.839
Data_train1	180	60	32	33	0.866	0.901	0.877	0.876
Data4_train	513	124	73	87	0.823	0.823	0.775	0.775
Data5_train	370	70	59	58	0.847	0.847	0.901	0.886
Data6_train	381	69	65	64	0.835	0.835	0.859	0.869
DLBCL	65	13	12	19	0.799	0.734	0.8	0.761
GDS2771	299	81	42	42	0.798	0.798	0.801	0.783
GDS2819_1	303	39	15	17	1	1	1	1
GDS2819_2	436	149	60	80	0.948	0.981	0.957	0.921
GDS2901	88	17	4	4	1	1	1	1
GDS2960	33	7	5	6	0.99	0.99	0.977	0.977
GDS2961	49	13	8	5	0.86	0.86	0.829	0.784
GDS2962	45	11	8	7	0.877	0.914	0.924	0.883
GDS3116	142	23	30	32	0.852	0.852	0.868	0.853
GDS3257	131	17	8	9	1	1	1	1
GDS3929	376	74	45	32	0.809	0.809	0.81	0.774
GDS4103	265	71	54	51	0.933	0.933	0.923	0.923
GDS4109	142	38	24	20	0.936	0.936	0.924	0.947
GDS4222	454	84	73	93	0.974	0.974	0.97	0.97
GDS4318	275	64	40	53	0.923	0.923	0.97	0.942
GDS4336	200	66	30	20	0.928	0.928	0.916	0.916
GDS4431	537	100	85	134	0.827	0.827	0.817	0.817
GDS4600	472	124	94	114	0.983	0.983	0.979	0.979
GDS4837_1	413	130	57	51	0.916	0.916	0.828	0.809
GDS4837_3	316	48	54	51	0.96	0.96	0.969	0.967
GDS4901	220	60	44	30	0.931	0.966	0.919	0.913
GDS4968_0	226	40	37	38	0.905	0.905	0.913	0.907
GDS4968_1	224	40	40	36	0.923	0.946	0.939	0.932
GDS5037_0	243	140	49	62	0.825	0.857	0.867	0.867
GDS5037_2	293	69	47	73	0.756	0.756	0.789	0.78
GDS5047	185	41	9	11	1	1	1	1
GDS5083	195	60	29	40	0.862	0.862	0.872	0.847
Leuk_3c0	34	5	7	7	0.989	0.989	0.986	0.986
Leuk_3c1	33	5	8	8	0.981	0.981	0.98	0.98
Ovarian	192	23	9	11	1	1	1	1
plySRBCT	17	3	0	1	1	1	1	1
prostate	93	34	16	15	0.919	0.932	0.927	0.921

6. Conclusion

In this paper, we presented two parallelization schemes for MeLiF algorithm, which search for the best combination of simple feature selection algorithms. Experiments showed that these two schemes, namely PQMeLiF and MAMeLiF, demonstrated linear speed improvement over the number of used cores comparing to single-threaded MeLiF without loss in classification quality. Furthermore, these algorithms sometimes showed better feature selection quality. This could be explained with the fact, that this methods search through more points than original MeLiF method. That happens due to delay in thread synchronization.

As our future research, we will try to make some estimations on search space split size in MAMeLiF method depending on the number of cores. Proper split can lead to better computational results. Also, we will apply more tests with different system configurations in order to find Amdahls optimum for each algorithm.

Acknowledgments

The research was supported by the Government of the Russian Federation (grant 074-U01) and the Russian Foundation for Basic Research (project no. 16-37-60115).

References

Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.

Benjamin Auffarth, Maite López, and Jesús Cerquides. Comparison of redundancy and relevance measures for feature selection in tissue classification of ct images. In *Industrial Conference on Data Mining*, pages 248–262. Springer, 2010.

Verónica Bolón-Canedo, Noelia Sánchez-Maroño, and Amparo Alonso-Betanzos. An ensemble of filters and classifiers for microarray data classification. *Pattern Recognition*, 45 (1):531–539, 2012.

Verónica Bolón-Canedo, Noelia Sánchez-Maroño, Amparo Alonso-Betanzos, José Manuel Benítez, and Francisco Herrera. A review of microarray datasets and applied feature selection methods. *Information Sciences*, 282:111–135, 2014.

Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: Stochastic and adversarial bandits. In *COLT*, pages 42–1, 2012.

Thomas Desautels, Andreas Krause, and Joel W Burdick. Parallelizing exploration-exploitation tradeoffs in gaussian process bandit optimization. *Journal of Machine Learning Research*, 15(1):3873–3923, 2014.

Ilya Isaev and Ivan Smetannikov. Melif+: Optimization of filter ensemble algorithm with parallel computing. In *IFIP International Conference on Artificial Intelligence Applications and Innovations*, pages 341–347. Springer, 2016.

Pooria Joulani, András György, and Csaba Szepesvári. Online learning under delayed feedback. In *ICML*, pages 1453–1461, 2013.

Ron Kohavi and George H John. Wrappers for feature subset selection. *Artificial intelligence*, 97(1):273–324, 1997.

Thomas Navin Lal, Olivier Chapelle, Jason Weston, and André Elisseeff. Embedded methods. In *Feature extraction*, pages 137–165. Springer, 2006.

Stuart Jonathan Russell and Peter Norvig. *Artificial intelligence: a modern approach*, volume 2. Prentice hall, 2009.

Yvan Saeys, Iñaki Inza, and Pedro Larrañaga. A review of feature selection techniques in bioinformatics. *bioinformatics*, 23(19):2507–2517, 2007.

Noelia Sánchez-Maroño, Amparo Alonso-Betanzos, and María Tombilla-Sanromán. Filter methods for feature selection–a comparative study. In *International Conference on Intelligent Data Engineering and Automated Learning*, pages 178–187. Springer, 2007.

Ivan Smetannikov and Andrey Filchenkov. Melif: Filter ensemble learning algorithm for gene selection. *Advanced Science Letters*, page (to be published), 2016.

Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press Cambridge, 1998.

This figure "spiral.png" is available in "png" format from:

http://arxiv.org/ps/1611.02047v1